

# **Should the Model for Security Be Game Theory Rather than Reliability Theory?**

Vicki Bier  
Center for Human Performance  
and Risk Analysis  
University of Wisconsin-Madison  
Madison, Wisconsin  
USA  
bier@engr.wisc.edu

## **Abstract**

Protecting against intentional attacks is fundamentally different from protecting against “acts of nature” or “accidents.” For example, an earthquake will not become stronger or “smarter” just because we have fortified our systems to protect against it. By contrast, an intelligent and determined adversary may adopt a different offensive strategy to circumvent (or destroy) our protective security measures. Therefore, a good defensive strategy must consider the adversary’s behaviour. We discuss the use of game theory combined with traditional reliability analysis to analyze and defend against security threats to networked (series/parallel) systems. The results of such work yield insights into the nature of optimal defensive investments in networked systems to obtain the best trade-off between the cost of the investments and the security of the resulting systems. In particular, we discuss how the optimal allocation of resources to defensive investments depends on features such as the structure of the system, the cost-effectiveness of the defensive investments, and also the attacker’s goals and constraints.

## **1. Introduction**

After the September 11, 2001, terrorist attacks on the World Trade Center and the Pentagon, there has been increased interest in strategies for analyzing and reducing the risk of intentional attacks. Protecting against intentional attacks is fundamentally different from protecting against accidents (the traditional domain of reliability analysis). The title of this paper is facetious; reliability analysis has an important role to play in security analysis, especially in analyzing and protecting against security threats to complex engineered systems. However, security analysis must go beyond merely reliability analysis. In particular, an intelligent and adaptable adversary may adopt a different offensive strategy to circumvent (or destroy) our protective security measures. Game theory (Dresher 1961) provides a way of taking this into account.

The fact that attackers can modify their strategies in response to our defensive investment suggests that defense will generally be more costly when the adversary knows the system defenses. In fact, Ravid (2002) argues that “investment in defensive measures, unlike investment in safety measures, saves a lower number of lives...than the apparent direct contribution of those measures.” Thus, security improvements may be less cost-effective than they might initially appear.

This paper discusses game theory models and results for defending reliability systems against attacks by knowledgeable and adaptable adversaries. We consider the problem both from the perspective of a game between an attacker and a defender, and also from the perspective of a game between multiple defenders.

## **2. Security as a game between an attacker and a defender**

Bier and Abhichandani (2003) consider the security of simple series and parallel systems. The results suggest that defending series systems against informed and determined attackers is a difficult challenge. If the attacker knows about the system's defenses, the defender's options for protecting the system are limited. In particular, the attacker's ability to respond to the defender's investments deprives the defender of the ability to allocate defensive investments according to their cost-effectiveness. Instead, in series systems, if potential attackers know about any defensive measures, defensive investments must essentially equalize the strength of all components in order to be beneficial (Dresher 1961).

This emphasizes the importance of redundancy as a defensive strategy. Essentially, redundancy reduces the flexibility available to the attacker in choice of targets (since the attacker must now disable multiple redundant components in order to disable a system), and increases the flexibility available to the defender (since the defender can now choose which of several redundant components to defend, based on the cost-effectiveness of doing so). Traditional reliability design considerations such as spatial separation and functional diversity are also important components of defensive strategy, to help ensure that attacks against redundant components are likely to succeed or fail more or less independently of each other (i.e., to ensure that redundant components cannot all be disabled by the same type of attack). Our results also support the idea that secrecy and even deception can be important strategies for improving security and/or reducing defensive investment costs, especially in series systems (where information about system defenses is valuable in allowing the attacker to optimize the choice of targets).

### **2.1 Weakest link models and extensions in series systems**

The fact that optimal defense of series systems requires equalizing the strength of all components makes this essentially a weakest link model—defensive investment must be allocated only to the weakest component(s) if it is not to be wasted. This is generally consistent with Brookings Institution recommendation (O'Hanlon et al. 2002) to defend only the most valuable assets. Our models go beyond the simple heuristic to also consider the success probabilities of attacks against various targets. This is important, since terrorists take the probability of success into account in choice of targets; see Woo (2003). Thus, even if one component is more valuable than another, it may not merit as much defensive investment if the less valuable component has a higher probability of being successfully attacked.

However, real-world decision makers will generally want to hedge their bets by investing in additional components, to cover contingencies such as whether they guessed wrong about which targets are most attractive to attackers. Woo (2003) achieves this by assuming that attackers do not always target the component(s) that would yield the highest utility to them, and instead randomize their target selection, with the probability of selecting a particular target being a non-linear function of the target attractiveness. Ongoing work is attempting to find an optimal strategy that hedges its bets by assuming that attackers always target the most attractive component, but that defenders are uncertain about the attractiveness of the targets to attackers. Under this model, attackers will in general have different values for particular targets than defenders; for example, Al-Qaeda may prefer targets that are "recognizable in the Middle East" (Woo 2003). Moreover, defending one target can deflect attacks to targets that are less attractive a priori to attackers, but more damaging to defenders. Thus, Enders and Sandler (forthcoming) note that making hijackings more difficult can increase other types of hostage-taking; similarly, defending government facilities may make officials less secure when they venture outside those facilities.

## 2.2 Other extensions

It is clearly important in practice to extend the types of security models described above to more complicated system structures, including both parallel and series subsystems, and this is one place where reliability analysts can make a substantial contribution. For example, past work on least-cost diagnosis of reliability systems could be adapted to identify least-cost attack strategies, for use in game-theoretic models as a building block for identifying optimal (or near-optimal) defensive strategies for such systems.

Most models to date assume that the success probability of an attack on a component is a convex function of the resources invested in that component. This is often reasonable, but is still restrictive. For example, if some security improvements require a minimal level of investment, this would result in step changes in the success probability of an attack as a function of the level of defensive investment. Similarly, if security investment beyond some threshold deters potential attackers from attempting an attack, then the likelihood of a successful attack could decrease rapidly beyond that threshold. Such effects will result in the success probability of an attack being a non-convex function of the defensive investment (at least when the investment is not too large), making it more difficult to identify the optimal level of investment.

Finally, it would be worthwhile to extend our models to include the time dimension, rather than the current “snapshot” view of system security. This would allow us to model imperfect attacker information (including, for example, Bayesian updating of the probability that an attack will succeed based on a past history of successful and failed attacks), as well as the possibility of multiple attacks over time.

## 3. Security as a game between defenders

In addition to anticipating the effects of defensive actions on possible attackers, it also makes sense to consider their effects on the strategies adopted by other defenders. Some defensive actions (such as installation of car alarms) may increase risk to other victims. This type of situation can lead to overinvestment in security, because the payoff to any one individual from investing is greater than the net payoff to the entire society. Conversely, other types of defensive actions (such as vaccination or use of anti-virus protection software) decreases risk to other potential victims. This type of situation can lead to underinvestment in security, since potential victims may “free ride” on the investments of others.

Kunreuther and Heal (2003) consider a model of interdependent security where agents are vulnerable to “infection” from other agents. For example, consider the example of supply chain partners who share access to each other’s computerized enterprise management systems, insider their partners’ firewalls. In this context, not only will defensive investment on the part of one agent benefit other agents; agents may actually be unable to defend their own systems against infections spread (however unintentionally) by their partners, and may need to rely on their partners to protect them against such threats. Kunreuther and Heal assume that even a single successful attack can be catastrophic—in other words, that the consequences of a successful attack (e.g., business failure) are “so serious that it is difficult to imagine an alternative event with greater consequences.” In the context of this model, they show that failure of one agent to invest in security can make it unprofitable for other agents to invest in security, even when they would normally find it profitable to do so. Moreover, they show that this game can in some cases have multiple equilibrium solutions. In such cases, there is a role for coordinating mechanisms such as contracts or voluntary standards to ensure that players arrive at the socially optimal level of investment.

Ongoing work is extending these results to attacks occurring over time. In this model, differences in discount rates can lead some agents not to invest in security when it would otherwise be in their interests, if other agents choose not to invest. Differences in discount rates can arise due to participation in

industries with different rates of return, to risk of impending bankruptcy, to myopia. As in the static model, coordinating mechanisms can help to ensure that the socially optimal level of investment is achieved. Heterogeneous discount rates complicate the task of achieving security in an interdependent world, and an understanding of this phenomenon can thus be helpful in identifying promising solutions.

#### 4. Conclusions

Protecting engineered systems against intentional attacks is likely to require a combination of game theory and reliability analysis. Reliability analysis by itself will likely not be sufficient to address most security challenges, since it does not take into account the attacker's response to any reliability or security improvements. However, most current applications of game theory to security deal with individual components in isolation, and could benefit from reliability analysis tools to address risks to complex systems such as computer systems, electricity transmission systems, or transportation systems. To answer the question posed in the title, perhaps the right model for security analysis is to embed reliability analysis in a game-theoretic framework, making it possible to take advantage of the strengths of both approaches.

#### Acknowledgements

This material is based upon work supported in part by the U.S. Army Research Laboratory and the U.S. Army Research Office under grant number DAAD19-01-1-0502, the U.S. National Science Foundation under grant number DMI-0228204, the Department of Homeland Security under grant number EMW-2004-GR-0112, and the Midwest Regional University Transportation Center under project number 04-05. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author and do not necessarily reflect the views of the sponsors.

#### References

- Bier, V. M., and V. Abhichandani (2003). Optimal allocation of resources for defense of simple series and parallel systems from determined adversaries. *Risk-Based Decisionmaking in Water Resources* X, pp. 59-76. Reston, VA: American Society of Civil Engineers.
- Dresher, M. (1961). *Games of strategy*. Englewood Cliffs, NJ: Prentice-Hall.
- Enders, W., and T. Sandler (forthcoming). What do we know about the substitution effect in transnational terrorism? <http://www-rcf.usc.edu/~tsandler/substitution2ms.pdf>. In A. Silke and G. Ilardi (Eds.), *Researching Terrorism Trends, Achievements, Failures*. London: Frank Cass.
- Kunreuther, H., and G. Heal (2003). Interdependent security. *Journal of Risk and Uncertainty* 26 (2-3), 231-249.
- O'Hanlon, M., P. Orszag, I. Daalder, M. Destler, D. Gunter, R. Litan, and J. Steinberg (2002). *Protecting the American Homeland*. Washington, DC: Brookings Institution.
- Ravid, I. (2002). Theater ballistic missiles and asymmetric war. The Military Conflict Institute, [http://www.militaryconflict.org/TBM%20and%20Asymmetric%20War%20L2%20\(1\)](http://www.militaryconflict.org/TBM%20and%20Asymmetric%20War%20L2%20(1)).
- Woo, G. (2003). Insuring against Al-Qaeda. <http://www.nber.org/~confer/2003/insurance03/woo.pdf>. Insurance Project Workshop, National Bureau of Economic Research, Inc.